

**METHOD AND SYSTEM FOR RESEARCHING
PRODUCT DYNAMICS IN MARKET BASKETS IN
CONJUNCTION WITH AGGREGATE MARKET BASKET PROPERTIES**

Background of the Invention

Field of the Invention

The present invention relates to a method and system of data mining and, more particularly, to a method and system of data mining which determines product dynamics in market baskets in conjunction with aggregate market basket properties.

Description of the Related Art

Data mining is a well known technology used to discover patterns and relationships in data. Data mining involves the application of advanced statistical analysis and modeling techniques to the data to find useful patterns and relationships. The resulting patterns and relationships are used in many applications in business to guide business actions and to make predictions helpful in planning future business actions.

One of the types of data mining is called "association analysis," often referred to as "market basket analysis." Association analysis reveals patterns in the form of "association rules" or "affinities." An association rule between products A and B can be expressed symbolically as $A \rightarrow B$ which translates to the statement: "Whenever product A is in a market basket, then product B tends to be in the market basket as well." This is an

example of "product dynamics," i.e., the effect of the purchase of one product on another product.

In the folklore of data mining, one of the most repeated stories illustrating product dynamics is that of the alleged discovery that beer and diapers frequently appear together in a shopping basket. The explanation given in this tale is that when fathers are sent out on an errand to buy diapers, they often purchase a six pack of their favorite beer as a reward. Using the association rule discussed above, this example would be expressed as "diapers \rightarrow beer" or, translated, whenever diapers appear in a shopping basket, beer also tends to appear in that shopping basket.

There are a number of measures that have historically been used to characterize the importance of a particular association rule. In the context of market basket analysis, these measures are calculated in relation to all market baskets under consideration. The "confidence" of a rule " $A \rightarrow B$ " is the probability that if a basket contains A it will also contain B. The "support" of a rule is the frequency of occurrence of the rule in the set of all transactions. The "lift" of the rule is a measure of the predictive power of the premise A. Lift is a multiplier for the probability of B in the presence of A versus the probability of B without any prior knowledge of other items in the market basket.

For purposes of explanation, consider the following example: Table 1 illustrates ten typical transactions representing the market baskets for a given day at a small store. From the data in the table, it can be seen that diapers and beer appear together in some market

baskets and we can conclude that when a transaction contains diapers, there is a tendency for it to also contain beer. Diapers appear in six transactions (1, 3, 4, 8, 9, and 10) and beer appears in conjunction with diapers in four of these transactions (1, 3, 9, and 10). Therefore, the rule "diapers \rightarrow beer" has a confidence of $4/6 = 67\%$. Further, there are four of the ten transactions where beer and diapers appear together. This results in a value of $4/10 = 40\%$ for the support of the rule. Finally, beer appears in five of the ten transactions while it appears in four of the six transactions containing diapers. This means that if a basket was randomly chosen without any prior information about any of the transactions, there is a $5/10 = 50\%$ chance of finding beer. However, if we use the prior knowledge that if the basket contains diapers it has a good likelihood of also having beer, then the prospect of finding beer is improved if we choose from only baskets known to contain diapers, i.e., there is a $4/6 = 67\%$ chance of finding beer. Thus, the lift of the rule "diapers \rightarrow beer" is $67\%/50\% = 1.34$.

TABLE 1

TRANSACTION	MARKET BASKET
1	Diapers, beer, chips, soap
2	Chips, soap
3	Diapers, beer, soap
4	Diapers, chips, soap
5	Soap
6	Chips

TRANSACTION	MARKET BASKET
7	Beer, chips
8	Diapers
9	Diapers, beer, soap
10	Diapers, beer, chips, soap

5 Association analysis techniques discover all association rules that exceed set support and confidence thresholds. They also discover all sets of items that tend to occur in the same basket with a frequency that exceeds the support threshold; such sets are termed "frequent itemsets."

10 In recognition of the importance of data mining, tools have been developed to perform the various data mining and modeling techniques. One such tool is Intelligent Miner™ sold by IBM. Intelligent Miner has an outstanding algorithm for association analysis as part of its tool suite. Being general purpose tools, Intelligent Miner and other data mining tools for association analysis reach the point of inferring frequent itemsets and rules with their corresponding metrics of interest, such as support, confidence, and lift, but
15 go no further.

Association rules express facts deduced from data. They are true statements about the relationships observed in the data. These rules, along with their measures of confidence, support, and lift, can and should be used to generate theories or hypotheses

about the effects of future actions that change the conditions under which the original observations were made. These hypotheses need to be posed in the complex and dynamic retail environment where potentially thousands of stores and tens of thousands of items must be considered against the backdrop of pricing actions, promotions, campaigns, seasonality, and product availability. Furthermore, all actions and results should be measured against a matrix of revenue and profit rather than the abstract notions of support, confidence, and lift.

Existing tools for association analysis do not factor in information about advertising and promotion and thus do not assist in developing theories or hypotheses about their effects on product sales and product dynamics in market baskets. Moreover, association analysis as commonly employed focuses on product dynamics and does not analyze the aggregate properties of individual baskets; referred to herein as "market basket dynamics." If such analysis were conducted, it would yield data which would allow an understanding of overall buying behavior, measured at the level of market baskets, and what drives the overall buying behavior. Currently implemented association rules and frequent itemsets do not assist in determining information about the overall buying habits of the owner of a particular type of market basket; for example, what kinds of products would be found in a "high-gross margin" baskets or which products may drive such "high-gross margin" baskets.

Analysis of market basket data using data mining techniques, such as association analysis, is a recent development. Traditional methods for evaluating the effects of advertising and promotions on sales for a particular item of interest focus on aggregate financial measures. For example, traditional approaches would measure the overall value of shopping baskets that include or do not include the item of interest and compute how these measures change as a function of promotion-related actions. These methods do not consider the overall content of shopping baskets (i.e., they focus only on items of interest) and thus do not explain what these baskets tend to contain, nor do they reveal data which allows analysis of market basket dynamics. Without information about the relationships between the sale of various items and their promotion status it is not possible to explain any observed changes in the aggregate measures. Moreover, by lumping together all baskets that contain the item of interest to compute an aggregate value, these methods do not allow for the possibility of having various types of baskets all containing the item of interest but with different dynamics and thus different aggregate values.

Accordingly, a need exists for a method and system for utilizing data mining techniques to obtain and analyze data which allows individual market baskets to be characterized based on all of the items in the basket, so that, for example, an understanding of the purchasing habits of persons who possess baskets having such characteristics can be gained.

10.

5

15

20

analyzing the group sales data based on the market basket groupings and determining if any of the market basket groups display characteristics identified by the analysis parameters; and for all market basket groups which have been determined to display the characteristics, enhancing the market basket group data by embedding therein an "imaginary item" which identifies the characteristic(s) displayed by each market basket group. In a preferred embodiment, association analysis is performed on the enhanced market basket group data to generate association rules and frequent itemsets, and the association rules and frequent itemsets are displayed and archived.

The present invention will now be described with reference to the following drawings, in which like reference numbers denote the same elements throughout.

Brief Description of the Drawings

Figure 1 is a block diagram of the functional components of a system constructed in accordance with the present invention;

Figure 2 is a high level flowchart illustrating the overall process of the present invention;

Figure 3 is a flowchart illustrating the aggregate property enhancement process block of Figure 2;

Figure 4 is a flowchart illustrating the advertising/promotional enhancement block of Figure 2;

Figure 5 illustrates a three-level merchandising taxonomy;

Figure 6 illustrates a taxonomy relation table linking the category level to the SKU level of the taxonomy illustrated in Figure 5;

Figure 7 illustrates a taxonomy relation table linking the category level to the department level for the taxonomy illustrated in Figure 5;

Figure 8 illustrates a three-level advertising taxonomy;

Figure 9 illustrates a taxonomy relation table linking an item level to a flyer/page level of the advertising taxonomy illustrated in Figure 8;

Figure 10 illustrates a taxonomy relation table linking the flyer/page level to the flyer level of the advertising taxonomy illustrated in Figure 8; and

Figure 11 is a flowchart illustrating one example of the post-processing step of Figure 2.

Detailed Description of the Preferred Embodiments

Figure 1 illustrates an overview of the functional components of a system 100 for researching product dynamics in market baskets in conjunction with aggregate market basket properties in accordance with the present invention. While the examples given herein are directed to a standard retail environment, the present invention is not limited to such an application, and it is clear that the principles and methods of the present invention can find application in numerous other settings including electronic commerce ("E-

Commerce") over the Internet and any other application in which it is desirable to research product dynamics in market baskets.

As used herein, the term "Retailer" refers to a person or organization needing analysis of data and conclusions derived from the data. Examples of typical Retailers include the marketing department of a retail store or E-Commerce organization, a buyer for a retail store, or a website designer designing a website for an E-Commerce organization.

As used herein, the term "User" refers to a person, organization, or automated device that uses the present invention in connection with input from a Retailer to provide the desired data analysis and conclusions. Examples of typical users include a market analysis organization, a retailer who possesses a system in accordance with the present invention; or a web designer for an E-Commerce organization that has access to a system in accordance with the present invention.

As used herein, the term "Purchaser" refers to a customer of a Retailer who purchases items from the Retailer. Examples of typical Purchasers include an individual shopper at a retail store or an individual making purchases over the Internet from an E-Commerce organization.

The system 100 includes Retailer information input/output devices 102, an analysis server 104, a database server 106 and data input devices 108, 110. Retailer information input devices 102 provide means for inputting to the system information regarding the

needs and desires of a particular Retailer. The Retailer information can be manually input by the User of the system via, for example, a standard keyboard; in this scenario, the Retailer information is manually gathered by the User by interview, questionnaire, or other known information-gathering techniques. In a preferred embodiment, the Retailer directly
5 inputs the Retailer information by filling out an electronic questionnaire or interview form available on a website of the User, thereby transmitting the information to analysis server 104 via the Internet.

Analysis server 104 performs the association analysis on all data input thereto. The general operation of the analysis server 104 is based on the use of well known systems and
10 tools such as IBM's Intelligent Miner discussed above. As described more fully below, however, it is the utilization of the data mining tools, the enhancement of the data analyzed by the data mining tools and the post-processing and interpretation of the findings that represent some of the novel aspects of the present invention.

Database server 106 stores selected point-of-sale (POS) data, advertising and
15 promotion data, and product data, all of which is obtained from a variety of general data sources, including POS systems 108, marketing department databases 110 and the like.

Each of the elements of the system can communicate with each other in any known manner, for example, over a network connection or via standard cabling. Analysis server 104 includes an interface 112, a server 114 (e.g., an HTTP or Internet web server) and
20 a controller 116. Interface 112 allows the analysis server 104 and the other elements of

the system to communicate with each other in a known manner. Server 114 manages communications between the input devices 102 and the database server 106 in a known manner. Controller 116 controls the operation of server 114, including all communications between preparation engine 118, association analysis engine 120 and post-processing engine 122. In general, preparation engine 118, association analysis engine 120, and post-processing engine 122 are known components found in most data mining systems. However, the preparation engine 118 and post-processing engine 122 are utilized in a novel manner to achieve the results of the present invention, as described in more detail below.

Figure 2 is a high level flowchart illustrating the overall process of the present invention. At step 210 analysis and specification parameters are acquired. These parameters are input to the system via a Retailer input device (e.g., one of the Retailer input devices 102 of Fig. 1) and define what it is that the Retailer is interested in knowing about. This information could include, for example, details regarding which POS data the Retailer is interested in analyzing (store locations, specific time periods, product lines); which hierarchy to use (e.g., merchandise and/or advertisement taxonomy, described below); minimum support and confidence thresholds, and item constraints (e.g., which product item(s) to include or exclude from the analysis).

At step 220 the data required to perform the various analyses requested by the Retailer in step 210 is collected and prepared. The collection aspect of step 220 involves

the gathering of selected data from the general data sources, for example, the POS system databases 108 and marketing department databases 110, by the database server 106. This assures that only necessary information is utilized and unnecessary information is excluded.

5 The preparation aspect of step 220 involves the preparing of the collected data for association analysis by assigning identification numbers for each of the events or items under scrutiny, for example a particular market basket and the items in the market basket, and then inserting the data for the revenue resulting from the sale of each item; the cost of each item, and details of any advertising that may have been done for each item. The exact data inserted will depend on the analysis requested by the Retailer and may be very detailed or only very general in nature. All of the data to be inserted is available from the POS system database 108 and the marketing department database 110.

10 In accordance with the present invention, the data is also enhanced, as described further below, by (a) embedding the data with information regarding advertisements and promotions; and/or (b) by identifying the aggregate characteristics of each market basket and embedding the market basket data with information regarding these aggregate properties. The preparation step allows the association analysis step at block 230 to take into account and process the parameters requested by the Retailer and the enhancement of the data enables the present invention to provide the Retailer with additional relevant information that is not available using prior-art market basket analysis methods.

15

20

At step 230, an association analysis is performed in a known manner using standard association analysis algorithms to process the data that has been collected and enhanced in step 220. In a well-known manner, the association analysis step generates association rules for the data. However, as discussed below, the rules are much more useful because of the enhancements introduced in step 220. Thus, the post-processing steps performed on the enhanced data at step 240 (described in more detail below with respect to Fig. 11) involve processing of the data based on parameters not considered by prior art systems, and thereby yield significantly better information for presentation and archiving at step 250.

Finally, at step 260 a determination is made as to whether additional analysis is required or desired. If not, the process ends. However, if more analysis is required, then the process returns to step 210 and begins again. There are numerous situations when additional analysis might be required. For example, it may be desired by the Retailer to run an analysis before a promotion is implemented, during the promotion, and after the promotion has ended so that changes in purchasing behavior can be identified and evaluated; to compare different stores or groups of stores (e.g., on a regional level); or to run an analysis of products at both a category level and a department level.

Figure 3 is a flowchart illustrating what is referred to herein as the "aggregate property" enhancement process of block 220 of Figure 2. The aggregate-property enhancement process illustrated with respect to Figure 3 enables the discovery of patterns that characterize or discriminate market baskets having particular overall properties, i.e.,

the market basket dynamics. Aggregate properties of a market basket would be, for example, a market basket that has an overall negative gross margin or a market basket that has an overall "high" gross margin. As shown in Fig. 3, a logical grouping of data (e.g., all the items contained in a market basket) are identified as possessing, as a whole, one or more specified properties.

Referring now to Figure 3, at step 310, data pertaining to a particular market basket (e.g., the market basket of Purchaser "David" on August 30th) is input and, at step 312, a determination is made as to whether or not a specified property about the market basket is true. Thus, for example, if the specified property being analyzed is the gross margin of the entire basket, and the Retailer has determined that a "high" gross margin market basket would be any basket that has an overall gross margin exceeding \$50, at step 312, if the market basket input at step 310 has a gross margin that is \$50 or less, the process proceeds to step 316 and the ordinary "non-enhanced" market basket information is written to the analysis file, and at step 318, a determination is made as to whether or not there is another market basket to be analyzed. If there are no additional market baskets to be analyzed, the process ends; if there is another market basket to be analyzed, the process proceeds back to step 310 and the process is continues until all baskets have been analyzed.

If, at step 312, it is determined that the property "gross margin" of the market basket "David" is high, i.e., that it exceeds \$50, then at step 314 a designation indicating the existence of this property "added" to the basket to identify this property as a

characteristic of the basket. These designations are referred to as "imaginary items" and enhance the market basket data by categorizing the market basket as possessing the designated property. As an example, at step 314, an imaginary item "HM" is "added" to the basket (i.e., the market basket data is modified to include a designation "HM") to indicate that the basket is a high margin basket. The addition of the imaginary item can comprise a simple coding process, wherein an identifier is added to the data for the market basket "David" which, in this example, identifies the basket as a high margin basket. Each imaginary item type requested by the retailer must have a different code so that they can be distinguished from each other, and it is also desirable, for efficiency, to make the imaginary items easily distinguishable from real items.

The process then proceeds to step 316 where the market basket information, including the now enhanced basket information, is written to the analysis file. The process may be repeated for as many properties as desired so that a basket may possess plural imaginary items identifying plural aggregate properties of the basket.

As noted above, in addition to characterizing market baskets, the present invention also enables analysis of the effects of advertising/promotion on the sale of products. During the advertising/promotional enhancement process all items in the database are classified using a standard merchandise taxonomy; at the same time, however, the items are also classified using an advertisement taxonomy, which enables association analysis to

produce patterns that involve elements of the advertising media used. These taxonomies are described in more detail below with respect to Figs. 5-10.

Figure 4 is a flowchart illustrating what is referred to herein as the "advertising/promotional enhancement" process of block 220 of Fig. 2. This process enables association analysis to take into account the advertising status for items at the time of sale. Referring now to Figure 4, at step 410 data pertaining to a next (or first) item in a sequence of transactions is obtained. This data comprises information regarding an individual purchase of an item in a market basket. At step 412, a determination is made as to whether or not the item was being advertised and/or promoted when it was purchased (based on data gathered during step 220 of Figure 2). If it is determined that the item was not advertised or promoted, at step 414 information is added to the data corresponding to the item to create an "enhanced item" that identifies the item as a non-advertised item, and the process proceeds back to step 410.

If, at step 412, it is determined that the item was advertised or promoted when it was purchased, then at step 416 the item is designated as being an advertised item, details of the advertisement are added to enhance the data pertaining to the item. These details may include low-level elements of the advertising media such as the particular quadrant of a page on which an advertisement appeared, high-level elements of the advertising media such as the year in which an advertisement appeared, or mid-level details falling between the low and high-level elements. The advertising elements are used in an

advertising taxonomy, described below with respect to Figs. 8-10, to allow the advertising information to be associated with other pertinent information.

Next, the process proceeds to step 418 to determine if the same product was advertised in another advertisement when it was purchased. This is done because many times a single product is advertised in multiple locations, by different methods, etc. If it is determined that there was another advertisement running for the same product when it was purchased, then it is flagged as such and the process proceeds back to step 416 so that the item may be designated as being advertised more than once and the details regarding the additional advertising/promotion programs can be added. This is repeated until it is determined that there are no more additional ads/promotions for the item, at which point the process proceeds back to step 410 and continues until all transactions/items have been processed.

In order to be able to use the enhanced data in association analysis, each item in the database is classified using a standard merchandise taxonomy, as well as an advertisement taxonomy. This classification process is described with reference to Figures 5-10. Referring now to Figure 5, a simple 3-level merchandising taxonomy is described. In practical application, the merchandise taxonomy could consist of many more levels, depending on the desired level of "resolution" of the data analysis and the organization of the business of the Retailer.

The idea behind any taxonomy classification is to establish links between various levels of classification so that "children" within the taxonomy can be linked to items of common "ancestry." For example, the merchandise taxonomy of Fig. 5 shows a three-level taxonomy populated with typical, basic sales data: a Department level; a Category level; and an Item or SKU (e.g., any product identification code used in the retail industry) level. Taking beverages as an example of a particular type of merchandise, at the Department level, the descriptions might be "alcoholic beverages," "non-alcoholic beverages"; at the Category level, the descriptions might be "beer," "wine," "cordials"; and at the Item level, the descriptions might be "Heineken six-pack 12 Oz. Bottles," "Corona six-pack 12 Oz. Cans," "Gallo Merlot 750 ml.," "Kendall Jackson Pinot Noir 750 ml."

To be able to use this information, relationships have to be established to link this information from the lowest level to the highest level. For example, Fig. 6 illustrates a taxonomy relation table (exemplary information illustrated only) linking the Category level to the Item level for the previously described example. As can be seen in Fig. 6, Item's "Heineken six-pack 12 Oz. Bottles" and "Corona six-pack 12 Oz. Cans" are each associated or linked to the category "Beer"; and Item's "Gallo Merlot" and "Kendall Jackson Pinot Noir" are associated with the category "Wine."

Figure 7 illustrates a taxonomy relation table linking the Category level to the Department level for the same example. As can be seen in Fig. 7, each of the categories

"Beer" and "Wine" are separately associated with the Department "Alcoholic Beverages."

Creation of these taxonomies is what enables association analysis to identify patterns that involve items at various levels, and this is especially important when low-level, less-frequently occurring items are involved. Such items do not allow patterns to be easily established on their own; with the taxonomies, however, these low level items may be considered at a higher level where patterns may be more easily established (e.g., if "Gallo Merlot" is infrequently purchased, advertising for the category "Wines" may be analyzed instead, since the taxonomy establishes the link to both the low-level and higher-level categories.

A novel aspect of the present invention is the use of association analysis to associate elements of the advertising media at various levels with each other, as well as with items in the merchandise taxonomy. Obviously this "cross-taxonomy association" is not limited to these two taxonomies and it is understood that other taxonomies could be associated with each other as well.

Figure 8 illustrates a simple 3-level advertising taxonomy. This taxonomy also comprises three levels populated with advertising/promotion data: a Flyer level; a Flyer/Page level; and an Item level. By using a level (Item level) that is also used in the merchandise taxonomy, a link is established between the two taxonomies so that correlations between the two taxonomies can be made. As with the merchandise taxonomy described above, the addition of more levels in the taxonomy will increase the resolution

of the analysis. Continuing with the same example, at the Flyer level, the description might be "September 23rd Philadelphia Inquirer Flyers"; "September 30th Philadelphia Inquirer Flyers";..., December 22nd Philadelphia Inquirer Flyers", etc.; for the Flyers/Page level, the descriptions might be "September 23rd Philadelphia Inquirer Flyer, Page 2"; December 22nd Burlington County Times Flyer, back page"; and for the Item level, the descriptions would be the same as for the merchandise taxonomy.

Figure 9 illustrates a taxonomy relation table linking an Item level to a "Flyer/Page" level of an advertising piece according to this example. The "Flyer/Page" level refers to a particular page of a particular advertising flyer on a specific date. Referring to Fig. 9, it can be seen that the Item "Heineken six-pack 12 Oz. Bottles" is associated with Flyer/Page "Philadelphia Inquirer, September 23rd/Front Page", which means that an advertisement for the Heineken Item of the example appeared on the front page of a flyer that was included in the September 23rd edition of the Philadelphia Inquirer. Similarly, the Corona Item of the example is associated with the same advertising page and date ("Philadelphia Inquirer, September 23rd/Front Page"); the Gallo Merlot Item is associated with the back page of a December 22nd flyer included with the Burlington County Times ("Burlington County Times/December 22nd/Back Page"); and the Kendall Jackson Item is associated with the middle insert of the December 22nd Burlington County Times flyer ("Burlington County Times/December 22nd/Middle Insert"). Note that the Heineken Item was also advertised on the back page of the September 23rd Philadelphia Inquirer.

Likewise, Fig. 10 illustrates a taxonomy relation table linking the "Flyer/Page" level of the taxonomy to the "Flyer" level (e.g., The front page of the September 23rd Philadelphia Inquirer flyer is associated with the September 23rd Philadelphia Inquirer flyer). Through this simple example, it can be seen that the advertising for a particular Item can be associated with a particular page or pages in an advertisement, and/or a particular advertising date for a particular publication.

The taxonomy relations that are illustrated in Figs. 5-10 set up the use of these taxonomies in performing the association analysis. These relationships are utilized, as described above, to enhance the market basket data prior to subjecting the data to association analysis. In a known manner, the Items in the merchandise taxonomy can be linked to advertisements (if any) in the advertising taxonomy.

Figure 11 is a flowchart illustrating one example of post-processing step 240 of Figure 2 and, in particular, a post-processing step for identifying patterns that characterize or discriminate market baskets with particular aggregate properties. As previously discussed, the association analysis step performed at step 230 of Figure 2 generates a series of rules which characterize each market basket, and many of these rules may have been enhanced by use of imaginary items, enhanced items, and/or taxonomies during the preparation of the data for processing. During the post-processing step, the analyzed, enhanced data is used to develop conclusions about the data.

Referring to Fig. 11, at step 1110, a determination is made as to whether or not a rule has been generated by the association analysis step for the item of interest, for example, for imaginary item HM. If no rule is found, this indicates that the analysis is complete and the process terminates at that point. However, if at step 1110 it is determined that the next rule to be processed involves the item of interest, then at step 1114, a determination is made as to whether or not the item of interest is part of the premise (e.g., before the arrow) of the rule or the consequent (e.g., after the arrow) of the rule. For example, if the item of interest is imaginary item HM, and if item HM is part of the premise of the rule (e.g., $HM \rightarrow A + B$, meaning "whenever a high margin basket occurs, it tends to include both A and B"), the process proceeds to step 1120 and the lift value, which was calculated during the association analysis step, is analyzed. If, for example, it is determined that the lift is much greater or much less than 1, then it is considered an "interesting" rule (i.e., of interest to the Retailer) and the process proceeds to step 1122 where the rule is deemed to **characterize** market baskets that have the property HM as being high margin baskets (and allows the inference that there is a high likelihood that such baskets contain A and B) and then the process proceeds back to step 1110. If the lift is found to be at or near 1, the rule is considered "uninteresting" and it is not designated as having any particular property, and the process proceeds back to step 1110.

If, on the other hand, at step 1114 it is determined that item HM is not part of the premise (i.e., that item HM is part of the consequent, e.g., $A+B \rightarrow HM$, meaning "whenever A and B occur together in a market basket, then it is a high margin basket"), then at step 1116, a determination is made as to whether the lift value is much greater or much less than 1. If the lift value is much greater or much less than 1, then at step 1118 the rule is deemed to **discriminate** market baskets that have the property HM from other baskets that do not have the property (and allows the inference that since the basket contains items A and B, there is a high likelihood that the basket is a high margin basket) and then the process proceeds back to step 1110. If, at step 1116, there is a determination made that the lift value is at or nearly 1, then the process just reverts back to step 1110.

Although the present invention is described in connection with marketing research, it is understood that the techniques and methods described herein can be applied to any type of research in which it is desired to characterize data groupings and/or analyze the effects of a particular parameter (e.g., other than advertising). Further, it is understood that the properties of the market baskets can include information other than financial information, for example, a market basket can be characterized as containing advertised and/or non-advertised items and this information can be used by researchers as well.

Although the present invention has been described with respect to a specific preferred embodiment thereof, various changes and modifications may be suggested to one

[illegible]